Solar Physics
DOI: 10.1007/•••••-•••-•••-••••-•

The Helioseismic and Magnetic Imager (HMI) Vector Magnetic Field Pipeline: Optimization of the Spectral Line Inversion Code

R. Centeno[1] · **J. Schou**[2,3] · **K. Hayashi**[2] · **A. Norton**[2] · **J.T. Hoeksema**[2] · **Y. Liu**[2] · **K.D. Leka**[4] · **G. Barnes**[4]

© Springer ••••

Abstract The Very Fast Inversion of the Stokes Vector (VFISV) is a Milne-Eddington spectral line inversion code used to determine the magnetic and thermodynamic parameters of the solar photosphere from observations of the Stokes vector in the 6173 Å Fe I line by the *Helioseismic and Magnetic Imager* (HMI) onboard the *Solar Dynamics Observatory* (SDO). We report on the modifications made to the original VFISV inversion code in order to optimize its operation within the HMI data pipeline and provide the smoothest solution in active regions. The changes either sped up the computation or reduced the frequency with which the algorithm failed to converge to a satisfactory solution. Additionally, coding bugs which were detected and fixed in the original VFISV release, are reported here.

1. Introduction: What is VFISV?

The Very Fast Inversion of the Stokes Vector (VFISV; Borrero *et al.*, 2011) is a spectral line inversion code tailored and optimized to invert the full disk spectropolarimetric data of the *Helioseismic and Magnetic Imager* (HMI) instrument (Scherrer *et al.*, 2012; Schou *et al.*, 2012) on board the *Solar Dynamics Observatory* (SDO; Pesnell, Thompson, and Chamberlin, 2012). HMI is a filtergram-type instrument that observes (with one camera) the Stokes I, Q, U, and V at six wavelength positions across the Fe I 6173 Å spectral line for the full disk of the Sun with 16 million pixels (4096×4096 CCD) every 135 s. All of the data pipeline procedures, and the spectral line inversion code in particular, have limited allowable runtimes in order to keep pace with the data flow rate and prevent a processing backlog.

[1] High Altitude Observatory (NCAR), 3080 Center Green Dr., Boulder CO 80301
[2] HEPL Solar Physics, Stanford University, CA 94305-4085
[3] Max Planck Institute for Solar System Research, Max-Planck-Str. 2, 37191 Katlenburg-Lindau, Germany
[4] NwRA, 3380 Mitchell Lane, Boulder, CO 80301

arXiv:1403.3677v1 [astro-ph.SR] 14 Mar 2014

In the forward problem, VFISV solves the radiative transfer equation (RTE) for polarized light using the Milne-Eddington (ME) approximation to generate a set of Stokes profiles from a given model atmosphere (Unno, 1956; Rachkovsky, 1962). The ME approximation assumes that all the parameters describing the atmosphere are constant along the line of sight (LOS) except for the source function, which varies linearly with optical depth. In addition, the generation of polarized radiation is formulated in the classical Zeeman effect regime. Traditionally, ME models applied to polarized RTE problems use up to eleven free parameters to describe the atmosphere in which the Stokes profiles are generated. There are five thermodynamical parameters: the line-to-continuum absorption ratio, η_0, the Doppler width, $\Delta\lambda_{\rm D}$, the damping parameter of the Voigt function, a, and the components of the linearized source function, B_0 and B_1, respectively. The magnetic field vector is described by three variables, *i.e.* the magnetic field strength, B, its inclination with respect to the LOS, γ, and its azimuth in the plane perpendicular to the line of sight, ψ (which is referenced to a column of pixels on the HMI CCD and increases counter-clockwise). There are also two kinematic parameters: the Doppler velocity, v, and the macroturbulent velocity, $v_{\rm mac}$. The former characterizes the macroscopic plasma speed, while the latter is generally used to model a combination of unresolved plasma velocity fields and instrument smearing effects - it is usually expressed in the form of a convolution of the Stokes profiles with a gaussian function of width $v_{\rm mac}$. VFISV makes explicit use of the measured HMI transmission filter profiles in the spectral line synthesis and hence does not need to use $v_{\rm mac}$ to account for the instrumental broadening, while the other thermodynamic parameters compensate for the unresolved velocities and any residual instrumental broadening. A standard additional geometrical parameter known as the filling factor, α, quantifies the fraction of light within any given pixel that originates from a magnetized atmosphere.

The optimization scheme of VFISV, based on a Levenberg-Marquardt (LM) minimization algorithm (see Press *et al.*, 1992), takes a set of observed Stokes profiles and finds the model parameters that best describe the atmosphere in which they were generated. It achieves this by performing a non-linear minimization of a merit function, χ^2, that measures the similarity between the observed and synthetic Stokes profiles in a least squares sense.

VFISV is just one module among many in the data pipeline for the HMI instrument. It operates on Level1.5 data (the `hmi_S.720s` data-series) from the 'vector' camera of HMI averaged every 720 s (Hoeksema *et al.*, 2014) and its output is fed to a disambiguation code (see Barnes *et al.*, 2014) that resolves the 180° azimuth ambiguity of the magnetic field vector. The original version of VFISV (Borrero *et al.*, 2011) was developed before the launch of SDO, which took place on 11 February 2010. The spectral line inversion code has been further optimized since HMI data became available online. The purpose of this paper is to describe the changes implemented in the code and its output since the launch of SDO.

The current version of the code is referred to as `fd10`, which began as a number (in this case 10) assigned to identify tests of various versions of the code on the full-disk (FD) data. The data series produced by VFISV `fd10` in the HMI

pipeline processing is called `hmi.ME_720s_fd10` and is available through JSOC[1] (Joint Science Operations Center). Its name reflects the fact that it is a product of the Milne-Eddington inversion (ME) every 720 s with the `fd10` version of the VFISV code.

This manuscript is organized as follows. Section 2 reports on the changes in the code that alter the χ^2-space and hence determine which families of solutions are preferred over others, whilst Section 3 describes all of the procedures and alterations that have a direct impact on the speed performance of the inversion but not on the final solution found by the algorithm. Section 4 sums up the improvements in speed and performance of the code and some bug fixes are reported in the Appendix.

2. Changes in the χ^2 Space

The changes implemented in the VFISV code since 2010 that are discussed in this section have an impact on the position of the minimum of the merit function, χ^2, that the algorithm is designed to minimize. Changes in the depth of the minima or the steepness of the function along a given parameter have an effect on the path that the algorithm takes to find a solution and on the solution itself. Subject to the limits set on parameters as described in Section 2.3 and Table 1, the objective of the code is always to find the global minimum.

2.1. Weighting of the Stokes Profiles

The inversion of the HMI Stokes vector is based on a non-linear least-squares minimization of a merit function, χ^2, that measures the difference between the observed and synthetic Stokes profiles (*i.e.* the goodness of the fit). This merit function is defined as:

$$\chi^2 = \frac{1}{F} \sum_S W_s^2 \sum_\lambda \frac{(\mathrm{OBS}_S(\lambda) - \mathrm{SYN}_S(\lambda, \mathbb{M}))^2}{\sigma_S^2} \tag{1}$$

where F is the number of degrees of freedom (*i.e.* the number of data points minus number of free parameters in the model), OBS and SYN refer to the observed and synthetic Stokes profiles, respectively, and $\mathbb{M}$ represents the model atmosphere. The index S denotes the 4 Stokes parameters and λ denotes the wavelength positions. σ_S represents the photon noise, which is a function of the intensity and the polarization state[2]. This is needed because a different number of filtergrams go into constructing the intensity and the polarization images.

[1] http://jsoc.stanford.edu/

[2] The expressions used to calculate the noise for the Stokes parameters come from accounting for the gain, the exposure time, the temporal averaging and the number of filtergrams that go into calculating each Stokes parameter. The latter results in different values of σ for the intensity and the polarization: $\sigma_I = 0.118 \cdot \sqrt{I_\mathrm{C}}$ and $\sigma_{\mathrm{Q,U,V}} = 0.204 \cdot \sqrt{I_\mathrm{C}}$. This arises as a consequence of the chosen polarization modulation scheme, described in Hoeksema *et al.* (2014)

However, the intrinsic wavelength dependence of σ_S is not accounted for in `fd10`, rendering it a function of the continuum intensity and the Stokes parameter only. The weights, W_S, are a set of four values (one for each Stokes parameter), whose purpose is to emphasize or de-emphasize the relative importance of a given Stokes parameter with respect to Stokes I in its contribution to χ^2.

The amplitude of the Stokes V signal is usually $10^{-1} - 10^{-2} \times I_{\rm C}$ (where $I_{\rm C}$ is the continuum intensity) while Stokes Q and U are typically one order of magnitude smaller. For this reason, the contribution of the Q, U, and V profiles to χ^2 is often almost negligible compared to that of Stokes I. However, Stokes Q, U, and V carry most of the information about the magnetic field. While Stokes I is only sensitive to the magnetic field strength in strong field regions, it is independent of azimuth and only slightly dependent on the inclination with respect to the LOS.

There are a number of factors that can negatively impact the quality of the fit of the Stokes vector in an inversion. Flatfielding errors or inaccuracies in our knowledge of the instruments's filter profiles could lead to systematic "missfits" of the Stokes parameters. In the case of HMI, the latter is likely to have a larger effect than the former. However, a more important source of error is that the physical model, $\mathbb{M}$, used to describe the atmosphere is not sufficiently realistic. Whatever the reason, an error in the fitting of the Stokes profiles will scale with their amplitude in the expression of χ^2, and because Stokes I is orders of magnitude bigger than Q, U, and V, its contribution will override all of the others. In other words, a slight missfit of Stokes I will always yield a non-zero difference of $\mathrm{OBS}_I(\lambda) - \mathrm{SYN}_I(\lambda, \mathbb{M})$ in the expression of χ^2. Due to the much larger amplitude of Stokes I compared to Stokes Q, U, and V, this term can easily mask equivalent differences for the polarization profiles (*i.e.* $\mathrm{OBS}_S(\lambda) - \mathrm{SYN}_S(\lambda, \mathbb{M})$, with $S = Q, U, V$). In order to de-emphasize Stokes I relative to Q, U, and V, we use a set of weights, determined empirically, that balances out the amplitude of the different Stokes signals in the expression of χ^2.

Because the expression of χ^2 acknowledges dependence of the photon noise on the Stokes parameter ($\sigma_I \neq \sigma_{Q,U,V}$ and $\sigma_Q = \sigma_U = \sigma_V$), it is useful to define the concept of *effective weights* as $W_S^{\rm eff} = W_S/\sigma_S$. The weight scheme currently adopted for VFISV is $W_S \approx$ [1,5,5,3.5] (or, equivalently, $W_S^{\rm eff} \approx$[1,3,3,2]) for [I,Q,U,V], respectively (note that these weights are then squared in Equation (1)). The effect of the weights is to change the shape of the χ^2 function and how steep the gradients along the different model parameters are.

The Sun hosts a range of magnetic activity, producing an array of polarization signatures in spectral lines. It is common to weight the Stokes profiles with a custom set of weights for each pixel, in order to balance out the amplitudes of the four Stokes profiles. In the original version of VFISV, Borrero *et al.* (2011) proposed $W_{Q,U,V} = 1$ and $W_I < 1$, where the weight for Stokes I was lowered more or less depending on the value of the continuum intensity for the pixel under consideration. However, for the sake of homogeneity, VFISV `fd10` has been altered so that the weighting scheme is the same for each and every pixel. Other instrument teams have implemented inversion codes with constant values

for W_S so that χ^2 values can be more easily compared. For instance, MERLIN[3] (the spectral line inversion code for the *Hinode*/SOT-SP data pipeline, which is direct heritage of the *Advanced Stokes Polarimeter* inversion code (Skumanich and Lites, 1987)), operates with constant values $W_S^2 = [1, 100, 100, 10]$ (or, equivalently, $W_S^{\rm eff} = [1, 10, 10, \sqrt{10}]$) for $[I,Q,U,V]$ (Lites *et al.*, 2007). This ratio for the weighting of the Stokes profiles is standard for high spectral resolution data from spectrograph instruments. However, the spectral smearing induced by the width of the HMI transmission filter profiles produces qualitatively different spectral line shapes and depths that call for custom W_S values. Figure 1 shows the differences in the magnetic field strength (top) and inclination with respect to the LOS (bottom), retrieved from the inversion of an AR observed by HMI using two different sets of weights, namely $W_S^{\rm eff} = [1, 10, 10, \sqrt{10}]$ in the left panels, and $W_S^{\rm eff} = [1, 3, 3, 2]$ in the right panels. The MERLIN weighting scheme (left panels) does not work for the HMI data, and the inhomogeneity in the retrieved magnetic field parameters becomes apparent in strong field regions. The chosen set of weights for the HMI data, $W_S^{\rm eff} = [1, 3, 3, 2]$, proved to yield the smoothest solution inside ARs.

The HMI vector team adopted the policy of uniform weights throughout the solar disk. The final choice of weights was selected to optimize the inversion results in active regions (ARs) with the aim of achieving smooth solutions for the magnetic field inside the umbrae and avoid discontinueties that would arise from changing weights. As a side effect, the weighting scheme is far from optimal for quiet Sun, where the adopted set of weights leads to artifacts in the solution for the inferred magnetic fields - it exacerbates the "horizontal field effect" in which the inversion of pixels that have near-to-pure noise polarization signals will deliver a seemingly strong purely transverse magnetic field vector (Borrero and Kobel, 2011).

2.2. Reducing the Number of Free Parameters

Inverting the HMI data is challenging due to a combination of moderate spectral resolution and very coarse spectral sampling. Since there are only six wavelength sampling positions across the 6173 Å Fe I spectral line, we have a relatively small number of observations with which to constrain a model with up to eleven free parameters. Observing a single spectral line instead of a line pair (such as the 6302 and 6301 Å Fe I line pair) means we further decrease our ability to constrain the physical system. Using two spectral lines, as *Hinode*/SOT-SP does, takes advantage of the different line formation physics (as evident in the Landé factor) to distinguish between parameters otherwise redundant. In general, there are degeneracies in the thermodynamical parameters of the model when carrying out the inversion. One approach to improving the performance of VFISV was to reduce the number of free parameters and thus decrease its degrees of freedom.

Early tests (Borrero *et al.* (2011), Figure 14) showed that it was possible to fix the damping to a constant value and to ignore the macroturbulence without

[3]https://www.csac.hao.ucar.edu/

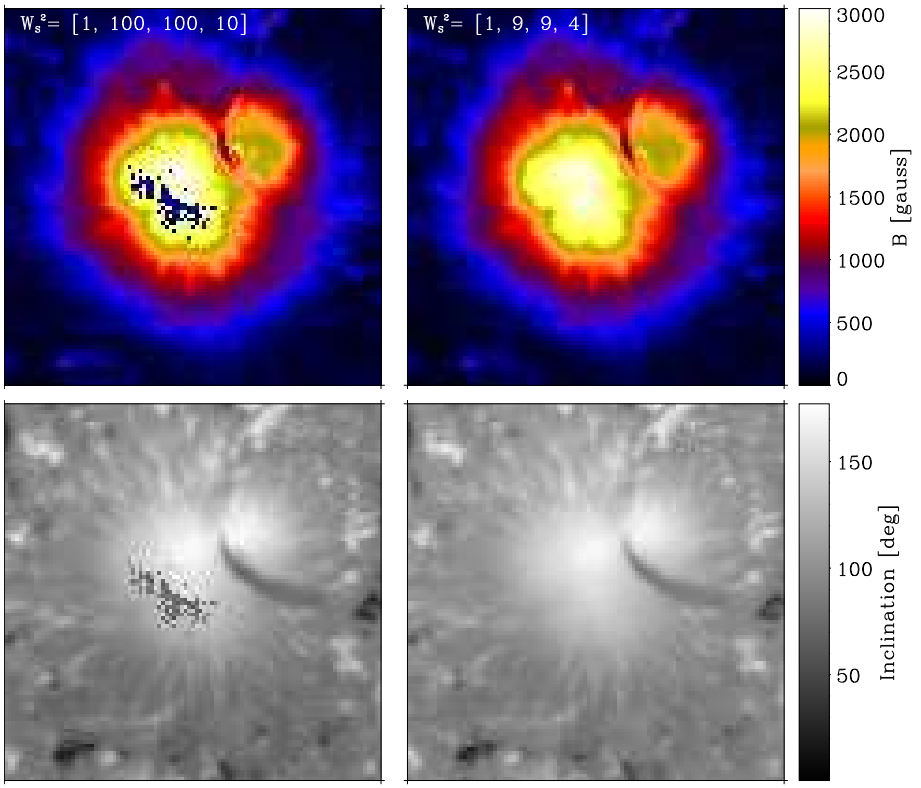

Figure 1. Magnetic field strength (top) and inclination (bottom) retrieved from the inversion of an AR using two different set of weights, that of the *Hinode*/SOT-SP MERLIN inversion code $W_S^2 = [1, 100, 100, 10]$ (left) and the one used in the fd10 version of the HMI data pipeline, $W_S^2 = [1, 9, 9, 4]$ (right). The inhomogeneity of the retrieved magnetic field in the umbra when using the MERLIN weights becomes apparent.

a loss in accuracy in the determination of the magnetic field vector due to the sparse sampling of the spectral line. Setting the damping to a fixed value reduced the computing time as well as the degeneracy of the results. Originally, the damping parameter was set to a constant value of $a = 1.0$, but further analysis showed that when solved for as a free parameter, the average value was closer to $a = 0.5$ for the height of formation of the 6173 Å line. Therefore, the damping parameter was set to a constant value of $a = 0.5$ for the production of `fd10` data. Furthermore, the magnetic flux, $\alpha \times B$, is well-constrained in the fit, but the individual values of α and B are not. For this reason, the filling factor is set to a constant value of $\alpha = 1$. A limitation introduced by assuming $\alpha = 1$ is that the returned magnetic field strength value is in fact an area-averaged quantity (Lites and Skumanich, 1990). In areas with unresolved magnetic structures such as plage, where averaged values of α are approximately 0.15 (Martinez Pillet, Lites

and Skumanich, 1997), the `fd10` field strength values will be significantly lower than the intrinsic field strength that would result from an inversion in which the filling factor were allowed to be less than one. In addition, the inclination angle will be biased towards the line of sight (Leka and Barnes, 2012).

2.3. Limits on Variables

Table 1. Upper and lower limits on the model parameters in VFISV.

Parameter	Symbol	Units	Lower limit	Upper limit
Doppler width	$\Delta\lambda_{\rm D}$	mÅ	1	500
Line-to-continuum ratio	η_0	dimensionless	1	1000
Doppler velocity	v	$\mathrm{cm\,s^{-1}}$	$-7\cdot10^5$	$7\cdot10^5$
Source function	B_0	$\mathrm{DN\,s^{-1}}$	$0.15\times I_C$	$1.2\times I_C$
Source function gradient	B_1	$\mathrm{DN\,s^{-1}}$	$0.15\times I_C$	$1.2\times I_C$
Field strength	B	gauss	5	5000
Inclination	γ	$^\circ$	0	180
Azimuth	ψ	$^\circ$	0	180

Occasionally, the iteration algorithm will find a location with a low χ^2 value, but unphysical parameters. In order to constrain the parameter search to a reasonable set of values, VFISV sets lower and upper limits to each of the physical parameters of the ME model atmosphere (see Table 1). These limits must be physically meaningful, at least in those cases where the ME parameter has a direct correspondence with a real physical magnitude. For instance, the magnetic field strength is not expected to exceed the 5000 G (gauss) limit, nor can it be less than 0 G. Even if an AR did harbor magnetic fields stronger than 5000 G, the dynamic range of the HMI instrument only reaches up to $\approx$ 3200 G when accounting for the large range of spacecraft velocity (Liu *et al.*, 2012), so the 5000 gauss (G) upper threshold covers the measurable range. Occasionally, when the observed Stokes profiles originate in the midst of very high field strengths (*i.e.* very dark places inside large sunspot umbrae) they challenge the dynamic range of the instrument and the inversion algorithm is unable to converge to the solution, hitting the hard limit set at 5000 G.

The limits on the LOS velocity, v, were determined taking into account typical plasma motions at the photosphere, the effects of solar rotation, the satellite's orbit around the Earth (which, alone, induces Doppler shifts of up to ±3.5 $\mathrm{km\,s^{-1}}$) and the Earth's orbit around the Sun. The inclination with respect to the LOS can only vary between 0 and 180°, and although the azimuth spans the entire 0-360° range, there is an inherent 180° ambiguity when determining it from the Stokes profiles alone. The terms of the source function should add up to the continuum intensity, thus, neither of them is allowed to exceed $1.2\times I_{\rm C}$. There is a strong degeneracy between these two parameters, hence an arbitrary lower limit was set, just to prevent negative values. The constraints on the rest of the variables were empirically derived and are justifiable from a χ^2 minimization perspective.

2.4. Regularization of χ^2

In the case of VFISV `fd10`, χ^2 represents a hypersurface of eight independent variables (*i.e.* the eight free parameters of the model atmosphere). In reality, due to limitations of the model and the data, these parameters are not all independent. Some degree of degeneracy exists among pairs and even combinations of several of them. This typically causes the χ^2 function to host narrow flat valleys where the LM algorithm cannot find the solution. One clear example of this is the degeneracy between the magnetic filling factor and the field strength in the weak field regime, *i.e.*, the Stokes profiles caused by a relatively-strong field with a small filling factor and those caused by a weak field with a large filling factor, are indistinguishable (see, for instance Martínez González *et al.*, 2006). In the `fd10` version of the HMI pipeline inversion code the filling factor is set to a fixed value of 1, as described in Section 2.2. Another example of degeneracy between parameters is described in Section 3.2.

A related problem is the presence of multiple minima in the χ^2 function. When there is a tendency for χ^2 to host two or more minima of similar depth in different parts of the parameter space, the algorithm will lead to each of these solutions with a certain degree of probability. Extensive tests of the inversion code on HMI data revealed a degree of degeneracy between η_0 and B. This particular degeneracy presented itself as a consistent double minima in the χ^2 function, with a bi-modal distribution of η_0 values. There would typically be two solutions with magnetic field strengths that differed by 100-300 G, one associated with a small value of η_0 and another one associated with a large value.

There is no physical reason to expect a bimodal distribution of η_0 values. The degeneracy between B and η_0 is largely a consequence of the limited spectral resolution of the HMI data. Milne-Eddington inversion results of *Hinode*/SOT-SP data (which have a spectral sampling of $\approx$ 21 mÅ per pixel) largely favor lower values of η_0. For this reason, high η_0 values were deemed unphysical, and an empirically determined regularization term was added to χ^2 in order to get rid of them:

$$\chi^2_{\rm new} = \chi^2_{\rm old} + \epsilon(\eta_0 - C)^2 \tag{2}$$

where $\epsilon = 0.002$ and $C = 5$. This term penalizes high values of η_0 in the merit function and reduces the number of pixels that are susceptible to a double minima behavior. Figure 2 shows the effects that different values of ϵ have on the histogram of η_0. The solid line corresponds to the case with no regularization term in χ^2, in which case a bi-modal distribution of η_0 becomes apparent (the arrow points to the secondary hump in the histogram, for $\eta_0 \approx 30$). The typical distribution of η_0 values for VFISV `fd10` is shown by the dashed line. With the constraint imposed by the regularization term, the secondary hump disappears, and the inversion favors lower values of η_0.

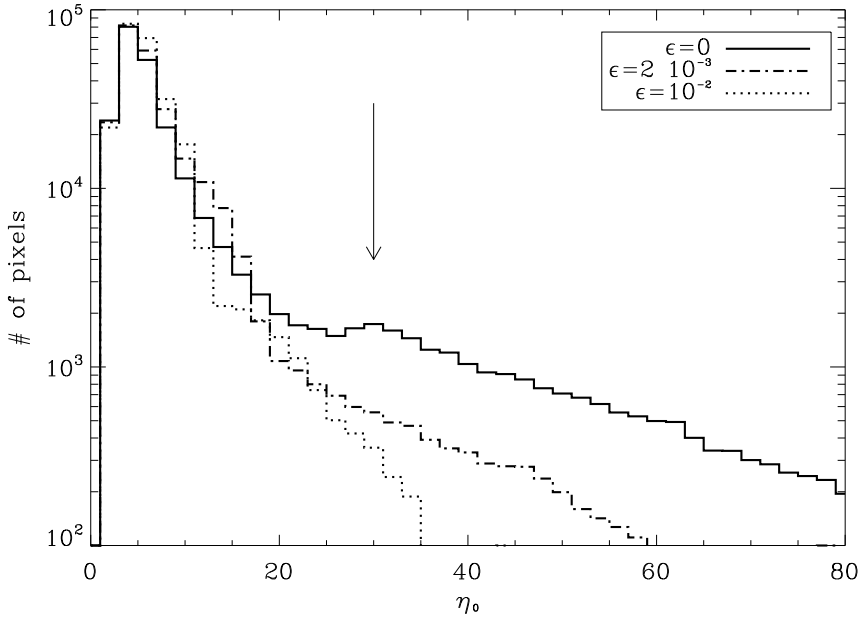

Figure 2. Effects of different regularization terms on the histogram of η_0 values. The solid line corresponds to no regularization, where the arrow points at the second hump in the bi-modal distribution of η_0. The dashed line corresponds to the regularization implemented in the VFISV fd10 pipeline (with $\epsilon = 2 \times 10^{-3}$), while the dotted line represents the histogram of η_0 for a larger value of $\epsilon = 10^{-2}$.

3. Speed Optimization

This section reports on modifications of the VFISV code aimed at speeding up the inversion without altering the shape of the merit function.

3.1. Explicit and Non-explicit Calculation of the Line Profile

Figure 3 shows the six HMI transmission filters as a function of wavelength (λ_0 corresponding to the central wavelength of the Fe I 6173 Å line). Each filter profile is characterized by a distinctive primary lobe and a number of smaller side lobes throughout the wavelength domain.

In its forward modeling module, VFISV synthesizes the Stokes profiles of the Fe I 6173 Å spectral line with a relatively high spectral resolution. It then integrates in wavelength (see Equation (3) in Borrero *et al.*, 2007) the products

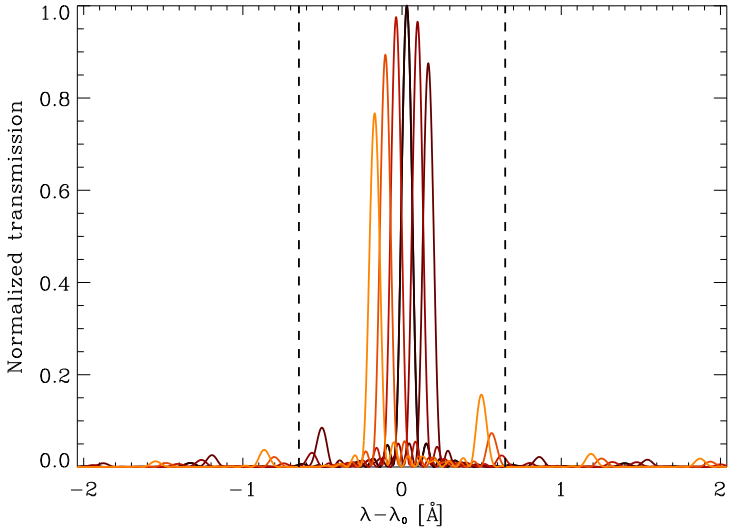

Figure 3. HMI transmission filter profiles as a function of wavelength, calculated for a range of [-2, 2] Å from the center of the HMI spectral line. The vertical dashed lines delimit the range in which the forward modeling of the spectral line is calculated explicitly in the inversion code.

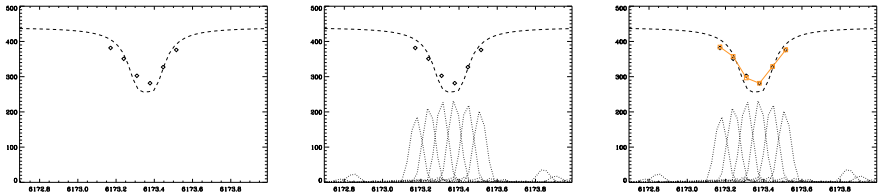

Figure 4. Sequence of images showing, schematically, how VFISV incorporates the HMI filter profiles into the spectral line modeling. From left to right: first panel shows the observed Stokes *I* profile (diamonds) and the synthetic spectral line (dashed). The second panel includes the HMI filter profiles for the observed pixel (dotted lines). The third panel superposes the synthetic filtered Stokes *I* profile (in orange), obtained as a result of integrating the synthetic spectral line under each of the HMI filter profiles.

of the spectral line with the six HMI filters to generate HMI-like Stokes profiles. The result of this process, schematically shown in Figure 4, is a set of synthetic filtered spectral profiles that are comparable to the data.

The wavelength range and sampling used to synthesize the HMI spectral line and the filter profiles are customizable parameters in VFISV. The larger the range and finer the sampling, the more accurate the modeling of the spectral line will be. When synthesizing the Fe I spectral line and applying the transmission profiles, it is necessary to account for the secondary lobes of the filters. Not doing so will lead to an underestimate, of up to 7%, of the amount of light that passes through the filters (see Couvidat *et al.*, 2012, for details on the calibration procedure for the HMI transmission profiles). Hence, it is important to extend the calculation as far into the wings of the line as practical. However, the wider

the spectral range, the larger the number of wavelength points in the spectral line synthesis calculation. This becomes rather computationally expensive. It is possible to reduce the number of points by lowering the spectral sampling of the synthesis, but this can only be done in detriment of the results.

A compromise has to be reached between increasing the spectral range and reducing the sampling. An alternative approach is to compute the contribution of the filter profiles far out in the continuum without doing the forward modeling of the spectral line in this region and without compromising the wavelength sampling. Figure 3 shows the HMI transmission profiles. The dashed lines represent the lower and upper wavelength boundaries within which the intensity profile of the HMI spectral line is always contained (in the absence of Zeeman splitting), when accounting for the Doppler shifts induced by the dynamic range of the satellite, the solar rotation, and typical photospheric motions[4]. In `fd10`, VFISV does the forward modeling only in this *inner* range (*i.e.* the spectral line is explicitly synthesized and the filter profiles are directly applied). In the *outer* range that goes farther into the wings the contribution of each filter profile is merely integrated, multiplied by the continuum intensity derived from the inversion and added to the synthetic Stokes I (this does not affect Stokes Q, U, and V because the continuum polarization is negligible). This is a reasonable approach when we work under the assumption that there is no spectral feature in the *outer* range. This strategy allows us to account for the contribution of the secondary and tertiary lobes of the instrument's transmission profiles without having to do the explicit calculation of the spectral line in this outer range.

A series of tests revealed that the extended spectral range should at least cover ± 2 Å from the line core, while the explicit calculation only needs to be carried out in the ± 0.65 Å range. Figure 5 shows the magnetic field retrieved from the inversion of an AR using the explicit forward modeling in the full ± 2Å spectral range. A comparison between this and the results of an inversion using only the *inner* ± 0.65 Å range, exposes the problems derived from not accounting for a large enough wavelength range in the synthesis of the spectral line (left panel of Figure 6). Here, the difference between the retrieved magnetic fields from both inversions is shown in the form of a scatter plot. Inside the sunspot (*i.e.* for large fields), the relative difference in the magnetic field strengths amounts as much as 5%. It is clear that the full wavelength range has to be considered in the calculation for the sake of accuracy. The right panel of Figure 6 shows the difference in the inferred magnetic field strengths between the case of the explicit full range calculation and the hybrid approach described above (in which the explicit calculation is done in the *inner* ± 0.65 Å range, and the remaining contribution of the spectrum is calculated implicitly). In this case, the number of pixels for which the difference in retrieved magnetic fields is large has been significantly reduced.

[4]Assuming $\pm 3.5\,\mathrm{km\,s^{-1}}$ line-of-sight velocity between the Sun and spacecraft, $\pm 2\,\mathrm{km\,s^{-1}}$ of solar rotation and $\pm 1.5\,\mathrm{km\,s^{-1}}$ of photospheric plasma motions, the possible induced Doppler shift can amount up to ± 0.144 Å. Under typical photospheric conditions, the spectral line is less than 1Å wide (*i.e.*, less than 0.5Å at each side of the position of the core), so the entire profile will always be contained in a ± 0.65 Å range around the central wavelength, in the absence of a magnetic field that would induce a Zeeman splitting.

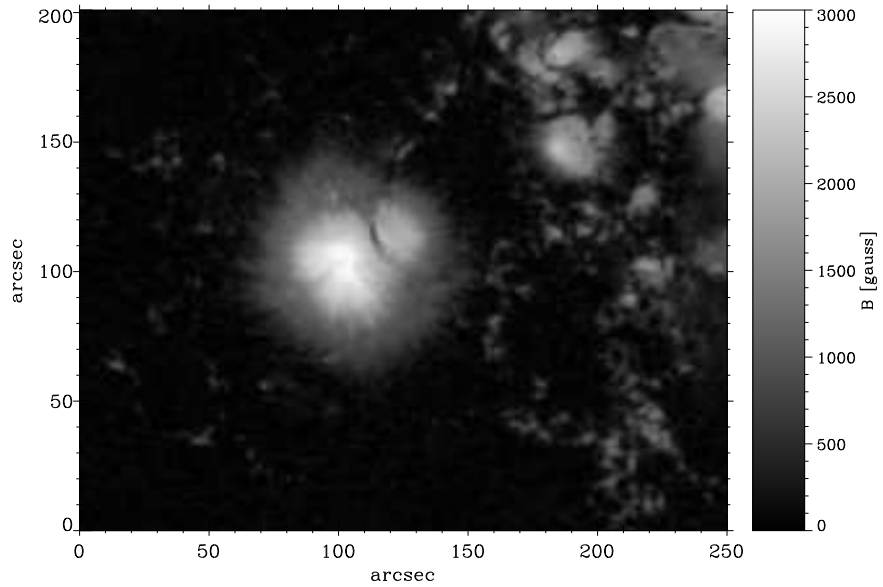

Figure 5. Magnetic field retrieved from an inversion where the full ±2 Å wavelength range is considered explicitly.

Since this hybrid approach allows us to limit the explicit calculation to a third of the total necessary spectral range, this technique results in a speed up of a full disk inversion by a factor of ≈ 2.75.

3.2. Variable Change

In exploring the details of the evolution of χ^2 during the inversion of selected pixels using the original VFISV code, it became apparent that the path to convergence was not always the fastest or the most efficient one. This stemmed, partly, from the degeneracy among the physical parameters in the model atmosphere. For instance, there is a strong coupling between the terms of the source function, B_0 and B_1. While the two parameters are difficult to determine individually, their sum is tightly constrained by the continuum intensity.

While this is not, in general, a major problem for an algorithm using the full Hessian, we nonetheless chose to fit for B_0 and $B_0 + B_1$, in order to make the inversion more well behaved.

When looking at the shape of the χ^2 surface close to convergence as a function of different pairs of variables, the degeneracy problem becomes obvious. The left

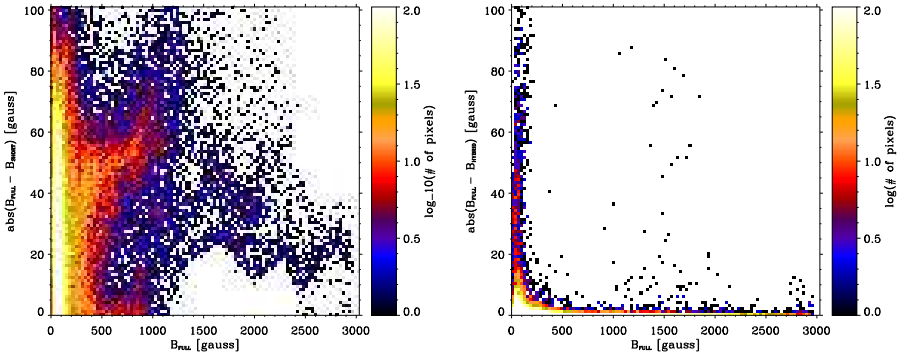

Figure 6. Effects of not including the secondary lobes of the HMI transmission profiles in the forward modeling. The left panel shows a density scatter plot of the difference between the magnetic field strength obtained from the full ±2Å explicit calculation and the result of the inversion using only the *inner* ±0.65 Å range. The right panel shows a scatter plot of the difference between the magnetic field of the full range explicit calculation and the result obtained using the hybrid approach. The differences in the retrieved magnetic field are visibly reduced when using the hybrid calculation.

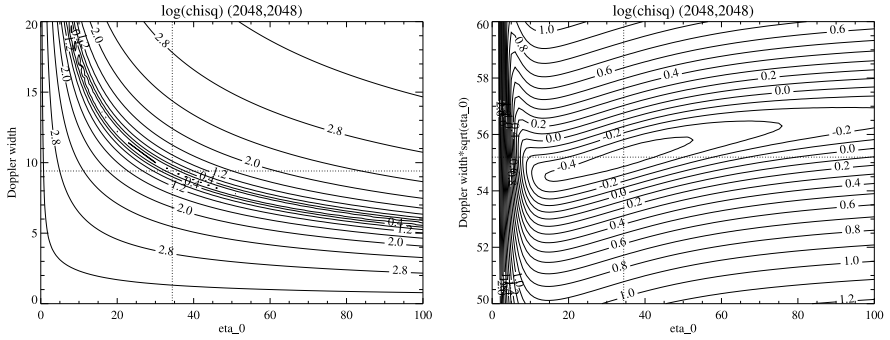

Figure 7. Contour plots of the χ^2 surface as a function of two parameters. Left: logarithm of χ^2 as a function of η_0 and $\Delta\lambda_D$ for the inversion of a the central pixel of the HMI CCD ([2048,2048] in the CCD coordinates). Right: exact same plot represented, instead, as a function of η_0 and $(\Delta\lambda_D \times \sqrt{\eta_0})$. The elliptical shapes of the χ^2-contours in the second case show that the change of variable will help the algorithm localize the solution (*i.e.* the minimum χ^2) faster.

panel of Figure 7 shows the logarithm of the χ^2 surface as a function of η_0 on the x-axis and the Doppler width, $\Delta\lambda_D$, on the y-axis. The contours show a narrow *curved valley* of minimum χ^2 values, along which different combinations of the two parameters result in a very similar goodness of the fit. Such a curved valley tends to lead to very poor performance of the LM algorithm. Implementing a change of variable, that makes the contours of the χ^2 surface closer to elliptical (see right panel of the same figure), leads to a χ^2 surface that is better approximated by a quadratic form in the parameters (that is, has a larger range

of validity of a Taylor expansion using the first and second partial derivatives) and, in turn, to a much improved convergence of the the LM algorithm.

Two changes of variable were implemented:

- B_0 and $B_1 \rightarrow B_0$ and $(B_0 + B_1)$
- η_0 and $\Delta\lambda_{\rm D} \rightarrow (\Delta\lambda_{\rm D} \times \sqrt{\eta_0})$ and $\Delta\lambda_{\rm D}$

The variable changes are computed just before the matrix inversion and undone right after obtaining the improved model parameters (see a schematic chart of the LM algorithm in Borrero *et al.* (2011)). This means that the forward modeling and the definition of χ^2 remain unchanged, and so does the physical model of the atmosphere. By definition, these variable changes do not alter the position of the minimum, expressed in the original variables, and so the position of the global minimum is unchanged. However, its path through the parameter space becomes more efficient after the variable change, which results in fewer iterations before convergence and hence a speed up of the inversion. The subsequent increase in speed achieved by the two variable changes described above results in a $\approx 10\%$ reduction in computing time for a full disk map in the HMI pipeline.

3.3. Iteration Algorithm and Testing

As part of the effort to speed up the inversions, the iteration algorithm of VFISV was significantly improved. Standard LM techniques are good at finding local minima in the minimization problem. However, they lack the capability to search the entire parameter space in pursuit of a global minimum. The approach that was followed for the `fd10` version of the HMI pipeline inversion has two levels of operation. The lower level takes an initial guess for the parameters and uses a LM algorithm to find a local minimum. The upper level tries to determine if the best solution found so far is the global minimum, and if not, a new initial guess is provided and the module calls the lower level again to try to find the global minimum; we call this a *restart*.

Several upper-level algorithms were tested against one another. In order to assess their performance, the quality of the fits obtained needed to be evaluated and compared to the time spent converging to the solution. The selected time criterion was the average number of iterations per pixel. While this is not exactly proportional to the running time, it is a good proxy. The criterion chosen to assess the quality of fits was the fraction of pixels for which the fitted χ^2 was within a given tolerance of that at the global minimum.

In order to evaluate the algorithms, a number of test datasets were selecthed that varied in complexity of the magnetic field and covered a range of line-of-sight spacecraft-to-Sun velocities. Unfortunately, the solution of the inversion problem (*i.e.* the global minimum) for real observations of the Sun is unknown. Therefore, the performance of the different algorithms had to be tested against some approximation to the solution.

In order to find the best estimate of the global minimum, a fixed version of VFISV was used to invert the test dataset one hundred times, each time with a different initial guess that was chosen at random spanning the entire parameter

space. For each inversion the code was forced to do two thousand iterations, allowing for restarts whenever the algorithm converged to a solution or failed to converge at all. For each pixel, the solution with the lowest value of χ^2 among the one hundred inversions was selected to represent the global minimum. A composite set of inversion results was put together using the best solution for each pixel in the field of view. This was labelled as the *gold standard* (GS) and has an associated $\chi^2_{\rm GS}$ map.

The lower level of the algorithm is a fairly standard LM algorithm. The control parameter λ is lowered by a factor of 5 if χ^2 decreases by more than a trivial amount (0.0001), with a lower bound for the value of 0.0001, and otherwise increased by an amount that depends on λ (10 if $\lambda \leq 0.01$ and 20 otherwise). At each step, the parameter limits are enforced, as are limits on the changes in some variables. It should also be noted that the model derivatives are only calculated when χ^2 improves (Borrero *et al.*, 2011), thereby saving a considerable time over always calculating them. The iteration stops if λ exceeds a certain value (100) or if the total number of iterations exceeds an upper bound (200).

At the end of each execution of the lower layer, the upper layer determines if a new initial guess for the model parameters should be tried, in which case, it generates one. A number of variations of the upper-level algorithm were thoroughly tested. They mostly differed in the way the new initial guess was chosen. The criteria for when a new guess and a subsequent restart is implemented were determined in a semi-empirical manner and were tailored to the `hmi.S_720s` data series. What follows is a description of the condition clauses for the upper-level algorithm implemented in `fd10` in order to find a global minimum in the inversion problem. This particular set of conditions was chosen because it showed the best compromise between speed and convergence to the *gold standard* [5].

At the end of the first execution of the lower level code, the algorithm will likely exit the iteration loop and the inversion will spit out a solution. However, for the pixels that meet one of the following criteria, one single restart will be enforced:

- The field is low ($B < 30\,$G) or high ($B > 300\,$G).
- The χ^2 is larger than 20.
- $\eta_0 < 5$ for $B < 300\,$G.
- Points with large magnetic field inclinations or low values of the Doppler width ($\Delta\lambda_{\rm D} < 15$).

At the end of the second execution, an upper-level decision to implement further restarts will be made based on the following criteria:

- Low ($B < (35 - N)\,$G) or high ($B > 500\,$G) fields , where N is the number of restarts already performed.
- Inclination within 30 degrees of vertical.
- Very low field ($B < 45\,$G) points with $\eta_0 < 3.5$ and $\Delta\lambda_{\rm D} > 30$.

[5]Note that the algorithm used to produce the *gold standard* was not a viable one for the HMI pipeline processing due to its high computation expense.

For the first and all subsequent odd restarts, a tailored guess is made as to a good new starting point in the parameter space. For the even numbered restarts, a random model perturbation is used. In most cases, the new initial guess will converge to the same solution as an earlier one. An additional control parameter is flagged when, after the first restart, the algorithm shows signs of heading inevitably to the same solution again, in which case it forces the exit of the iteration loop to avoid unnecessary iterations. The algorithm will never allow the total number of iterations (*i.e.* the sum of the iterations performed in all the restarts) exceed the maximum allowed (200).

The combination of the LM with this upper-level algorithm ensures that the global minimum is almost always found without using an excessive amount of computing time. The statistics for VFISV `fd10` code are as follows:

- Of all the pixels, 0.008% have not converged to within 0.1 of the $\chi^2_{\rm GS}$ value and 0.185% have not converged to within 0.01 of the $\chi^2_{\rm GS}$.
- Of the pixels with $B > 300$ G, 0.041% have not converged to within 0.1 and 0.097% not to within 0.01 of $\chi^2_{\rm GS}$.
- The rms error in B is 3.00 G.
- The average number of iterations is 28.22.

3.4. Parameter Initialization

The first initial guess of the atmospheric parameters for the inversion algorithm is mostly tailored to each pixel, although some of the parameters are given the same initial value for the entire field of view. This is true for the magnetic field strength, which is always initialized at 150 G, and its LOS inclination, which is started at 90°. However, other parameters are initialized differently for each individual pixel. A weak field approximation is used to calculate the initial value for the azimuth ($\psi = (1/2)\arctan(Q/U)$), and a rough guess, based on the wavelength shift of the spectral line, is estimated for the velocity. The source function and its gradient are constrained to add up to the continuum intensity, and an estimate of the Doppler width ($\Delta\lambda_{\rm D}$) is obtained from the shape of the spectral line. However, these first guesses are not crucial because the restart system described above erases any memory of the initialization, especially for the pixels with a relatively large magnetic signal, for which restarts are enforced until the maximum allowed number of iterations is reached.

4. Conclusion

Some of the changes to the VFISV code reported in this paper were aimed at speeding up the code while having little or no impact on the accuracy of the solution found by the inversion algorithm. These include measures that prevent the algorithm from performing unnecessary calculations and from following inefficient paths through the parameter space in pursuit of the solution. Certain modifications tried to palliate other performance issues that arise from the degeneracy among the atmospheric model parameters (which lead to multiple

minima in the χ^2) from the sometimes inadequate simplicity of the model chosen to represent the complexity of the real Sun, and from the data themselves, which have limited spectral and spatial resolution, are affected by photon noise, *etc.*

Throughout this manuscript we organized the changes into two classes: those that altered the shape of the χ^2 function (and hence the model atmosphere), and those that did not. However, they are not all independent of one another, and they operate together to improve the overall efficiency and performance of VFISV in the HMI data pipeline.

- A set of weights for the Stokes profiles that provided the smoothest possible solution inside active regions was chosen. This renders less than optimal results in weakly magnetized areas where the polarization profiles are heavily dominated by the noise. Because of the smearing effect of the transmission filter profiles of the HMI instrument, the standard weighting system used to invert datasets from spectrograph type of instruments (such as ASP or *Hinode*/SOT-SP), does not work with HMI data. A custom set of weights was chosen and applied uniformly over the entire FOV.
- A regularization term that penalizes high values of η_0 was added to the merit function. This mitigates the double minima problem in the parameter space and helps prevent the code from converging to unphysical solutions.
- Limits on the range of each atmospheric parameter were set to prevent the algorithm from probing extreme unphysical values.
- A diagnostic procedure that identifies possibly problematic pixels enforces a series of inversion restarts that attempt to localize the global minimum. Coupled to this, a discriminating convergence criteria ensures a very accurate solution while keeping the average number of iterations relatively low (≈ 30).
- A major speed improvement (about 64% reduction in computing time) was achieved, without compromising the accuracy of the solution, by considering a spectral range of ± 2 Å around the central wavelength but limiting the explicit forward modeling calculation to an inner ± 0.65 Å range. This hybrid approach accounts for the light that passes through the secondary lobes of the HMI transmission filter profiles without doing the detailed spectral line calculation far out into the continuum.
- An additional 10% increase in speed was obtained after performing two changes of variable on the atmospheric model parameters before solving the linear system of equations at each iterative step. This modification prevented the code from performing useless iterations in the parameter space by helping it find a more direct route to the solution.
- Coding errors were identified and corrected. Some of these are reported in the Appendix.

Acknowledgements Much of the reported effort was supported by Stanford University NASA Grant NAS5-02139 for SDO/HMI commissioning and pipeline code implementation; KDL and GB also acknowledge PO# NNG12PP28D/C# GS-23F-0197P from NASA/Goddard Space Flight Center. The National Center for Atmospheric Research is sponsored by the National Science Foundation.

Appendix

This section reports some coding mistakes that have been spotted in VFISV since its original release and that have been fixed in the `fd10` HMI data pipeline version. This version can be found through the JSOC (Joint Science Operations Center) CVS tree at Stanford.

A. Voigt Function

The Voigt function is a computationally intensive part of the forward modeling of the spectral line. In order to speed up the calculation, the actual Voigt function is only computed for small parts of the spectral range, close to the line core and far into the wings. In the intermediate wavelength regime the Voigt function is approximated by a Taylor expansion using a lookup table for the Taylor coefficients as a function of the central wavelength of the Voigt profile. Two bugs were found in the coding of the Voigt function of the spectral line.

An error in the wavelength indexing for the lookup table (inside the `voigt_taylor.f90` routine) resulted in sudden unphysical changes in the integrated area of the Voigt profile as a function of the central wavelength. This led to a periodic roughness pattern of the χ^2-surface that resulted in the appearance of multiple local minima along the magnetic field direction when the algorithm was approaching the solution. Instead of being evenly spread (due to the effect of the photon noise), the solutions would preferentially fall along ridges with an 11 gauss periodicity in the magnetic field.

A second coding mishap in the calculation of the Voigt profile was a missing factor of 2 in the computation of the imaginary part of the Taylor expansion of the Voigt function, which affected the magneto-optical effects in strong field regions. This error is heritage of the same missing factor of 2 in a FORTRAN routine (`voigt.f`, originally developed by Harvey and Nordlund), that calculates the Voigt function. This routine has been used in a number of spectral line inversion codes for the calculation of the line profiles. Albeit small, the effect on the synthetic Stokes profiles lead to inhomogeneous patterns in the inverted magnetic fields inside sunspot umbrae.

B. Uncertainties

The photon noise in the observed Stokes profiles propagates to the atmospheric model inferred by the inversion. When the algorithm has converged and the optimal model has been found, VFISV computes the formal uncertainties corresponding to each physical parameter. The variance associated with the retrieved model parameter a_j is taken to be proportional to the corresponding diagonal element of the inverse of the curvature matrix α_{jj}^{-1} (see, for instance, Bevington, 1969):

$$\sigma_{a_j}^2 = \frac{\chi^2}{N_{\mathrm{FREE}}} \alpha_{jj}^{-1} \tag{3}$$

where the curvature matrix is the second cross-partial derivative of χ^2 with respect to combinations of two of the model parameters:

$$\alpha_{jk} = \frac{1}{2}\frac{\partial^2 \chi^2}{\partial a_j \partial a_k}. \tag{4}$$

The non-diagonal elements of the inverse of the curvature matrix, α_{jk}^{-1}, are proportional to the covariances, σ_{jk}^2, between each pair, a_j and a_k. These elements give a sense of the degeneracy between any two given parameters of the model atmosphere.

An index switch inside the nested loop that calculates the elements of the covariance matrix (subroutine `get_err` of the `inv_utils.f90` module) led to huge uncertainty values that were many orders of magnitude too large. After correcting for this problem, the magnitudes of these formal uncertainties lie within reasonable values and compare to those obtained through a Monte Carlo experiment.

References

G. Barnes, K.D. Leka, A.D. Crouch, X. Sun, E.D. Wagner, J.Schou: 2014, *Solar Phys.* in preparation.

Bevington, P. R., 1969, *Data Reduction and Error Analysis for the Physical Sciences*, McGraw-Hill, New York, 164.

Borrero, J.M., Tomczyk, S., Norton, A.A., Darnell, T., Schou, J., Scherrer, P., Bush, R., Liu, Y.: 2007, *Solar Phys.* **240**, 177.

Borrero, J.M., Tomczyk, S., Kubo, M., Socas-Navarro, H., Schou, J., Couvidat, S., Bogart, R.: 2011, *Solar Phys.* **273**, 267.

Borrero, J.M., Kobel, P.: 2011, *Astron. Astrophys.* **527**, A29.

Couvidat, S., Schou, J., Shine, R., Bush, R.I., Miles, J.W., Scherrer, P.H., Rairden, R.L., 2012, *Solar Phys.* **275**, 285.

Hoeksema, J.T., Liu, Y., Hayashi. K., Sun, X., Schou, J., Couvidat, S., Norton, A., Bobra, M., Centeno, R., Leka, K.D., Barnes, G. and Turmon, M.: 2014, *Solar Phys.*, in press.

Leka, K.D., Barnes, G.: 2012, *Solar Phys.* **277**, 89.

Lites, B.W., Casini, R., Garcia, J., Socas-Navarro, H.: 2007, *Mem. Soc. Astron. Ital.* **78**, 148.

Lites, B.W., Skumanich, A.: 1990, *Astrophys. J.* **348**, 747.

Liu, Y., Hoeksema, J.T., Scherrer, P.H., Schou, J., Couvidat, S., Bush, R.I., Duvall, T.L., Hayashi, K., Sun, X., Zhao, X. : 2012, *Solar Phys.* **279**, 295.

Martínez González, M.J., Collados, M., Ruíz Cobo, B.: 2006 *Astron. Astrophys.* **456**, 1159.

Martínez Pillet, V., Lites, B. W., Skumanich, A.: 1997, *Astrophys. J.* **474**, 810.

Pesnell, W.D., Thompson, B.J., Chamberlin, P.C.: 2012, *Solar Phys.* **275**, 3.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., : 1992, Numerical Recipes in C, Second Edition, Cambridge University Press, Cambridge, 678.

Rachkovsky, D.N., 1962 *Izv. Krymskoi Astrofiz. Obs.* **28**, 259.

Scherrer, P.H., Schou, J., Bush, R.I., Kosovichev, A.G., Bogard, R.S., Hoeksema, J.T., *et al.*: 2012, *Solar Phys.* **275**, 207S.

Schou, J., Scherrer, P.H., Bush, R.I., Wachter, R., Couvidat, S., Rabello-Soares, M.C., *et al.*; 2012, *Solar Phys.* **275**, 229.

Skumanich, A., Lites, B.W.: 1987, *Astrophys. J.* **322**, 473S.

Unno, W, 1956, *Publ. Astron. Soc. Japan* **8**, 108.